Designing and using prior knowledge for phase retrieval

Eliyahu Osherovich*, Michael Zibulevsky, and Irad Yavneh

Computer Science Department, Technion — Israel Institute of Technology, Haifa, 32000, Israel

*Corresponding author: oeli@cs.technion.ac.il

In this work we develop an algorithm for signal reconstruction from the magnitude of its Fourier transform in a situation where some (non-zero) parts of the sought signal are known. Although our method does not assume that the known part comprises the boundary of the sought signal, this is often the case in microscopy: a specimen is placed inside a known mask, which can be thought of as a known light source that surrounds the unknown signal. Therefore, in the past, several algorithms were suggested that solve the phase retrieval problem assuming known boundary values [4,5,7,8]. Unlike our method, these methods do rely on the fact that the known part is on the boundary.

Besides the reconstruction method we give an explanation of the phenomena observed in previous work: the reconstruction is much faster when there is more energy concentrated in the known part [2, 5]. Quite surprisingly, this can be explained using our previous results on phase retrieval with approximately known Fourier phase.

1. Introduction

The phase retrieval problem amounts to signal reconstruction from the magnitude of its Fourier transform. The problems arises in several fields of physics including crystallography, astronomy, and certain branches of optics and microscopy. As with any inverse problem where the sought signal must be reconstructed from incomplete data, the main questions are: uniqueness of the reconstruction and design of an efficient reconstruction method that would be insensitive to possible errors in the measurements. It it trivial to show that, without additional information about the signal, the Fourier magnitude alone cannot provide sufficient information for successful reconstruction. Inasmuch as each Fourier domain component can

1

be assigned an arbitrary phase to accompany the measured magnitude resulting in a signal that fully corresponds to the measurement, namely the Fourier magnitude. One of the most widely used assumptions is that the sought signal $z(m, n)$ has limited spacial extent, i.e., one can assume that outside the rectangle $[0, M - 1] \times [0, N - 1]$ the signal values are zeros. In this case, it has been shown that in most cases the phase retrieval problem has a unique solution up to a trivial transformation. Possible variations being: constant phase factor, shifts, and axis reversal [1, 6]. Hence, in what follows we shall not treat in depth the problem of uniqueness and will concentrate on the efficient reconstruction method that provides good reconstruction results even in the case of nosy measurements.

In this paper we develop an algorithm for signal reconstruction from the magnitude of its Fourier transform in a situation where some (non-zero) parts of the sought signal are known. Although our method does not assume that the known part comprises the boundary of the sought signal, this is often the case in microscopy: a specimen is placed inside a known mask, which can be thought of as a known light source that surrounds the unknown signal. Therefore, in the past, several algorithms were suggested that solve the phase retrieval problem assuming known boundary values [4, 5, 7, 8]. Unlike our method, these methods do rely on the fact that the known part in on the boundary.

Besides the reconstruction method we give an explanation of the phenomena observed in previous work: the reconstruction is much faster when there is more energy concentrated in the known part [2, 5]. Quite surprisingly, this can be explained using our previous results on phase retrieval with approximately known Fourier phase [9].

2.   Review of existing methods

Before we proceed further it is important to note that the standard Hybrid Input-Output method [3] can be applied because the Fourier magnitude of the sought signal is known. However, this approach is not optimal as it does not use the additional information available in this case. Furthermore, HIO fails in case of complex-valued objects without tight support information.

Let us start with the notation used throughout the paper. As we already mentioned, the sought signal is denoted by $z(m, n)$. The signal is assumed to have a finite support, specifically, it vanishes outside the box $[0, M - 1] \times [0, N - 1]$, and our goal is, as before, to reconstruct it from the (oversampled) magnitude of its Fourier transform $|\hat{z}(p, q)|$, where

$$\hat{z}(p, q) = |\hat{z}(p, q)| \exp(i\phi(p, q)) \equiv \mathcal{F}[z(m, n)] . \tag{1}$$

In our implementation, F the *unitary* Fourier transform. This regularization constant is chosen with only one purpose: to make the distance (norm) in the Fourier domain equal that in the object domain. Thus, by the discrepancy (error) in the Fourier domain one can easily

estimate the error in the object domain. Hence,

$$\mathcal{F}[z(m,n)] = (PQ)^{-1/2} \sum_{m=0}^{P-1} \sum_{n=0}^{Q-1} z(m,n) \exp[-i2\pi(mp/P + nq/Q)] \,, \tag{2}$$

where $m, p = 0, 1, \ldots, P-1$ and $n, q = 0, 1, \ldots, Q-1$. In addition, we require the "two-fold oversampling" in the Fourier domain, that is $P = 2M - 2$, $Q = 2N - 2$. The purpose of this oversampling is to capture the information that $z(m,n)$ has only limited support. With these definitions, there exists one well-known relation between the squared magnitude of the Fourier transform and the linear (as opposed to cyclic) auto-correlation (denoted by $\star$) function of $z(m,n)$:

$$
\begin{aligned}
r(i,j) &\equiv z(m,n) \star z(m,n) \\
&\equiv \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} z(m,n) z^*(m-i, n-j) \\
&= \mathcal{F}^{-1}(|\hat{z}(p,q)|^2) \,,
\end{aligned}
\tag{3}
$$

where $r(i,j)$ is of size $(2M-1) \times (2N-1)$ pixels defined over the region $1 - M \leq i \leq M - 1$, $1 - N \leq j \leq N - 1$. This relation between the magnitude of the Fourier transform and the auto-correlation function of the sought signal was used by Hayes and Quatieri to develop an elegant algorithm for finding x (note that x is not necessarily equal to z due to possible non-uniqueness) when its boundaries are known [7,8]. The algorithm assumes that the boundaries, that is, the first row $z(0,:)$,[1] the last row $z(M-1,:)$, as well as the first and the last columns: $z(:,1)$, and $z(:,N-1)$, are known. The algorithm is iterative and at each iteration it recovers two new unknown rows (or columns) of x. The authors demonstrated that the k-th iteration of the algorithm reduces to a simple matrix (pseudo) inverse to solve the following system of equations

$$
\left[\, F \mid L \,\right]
\begin{bmatrix}
x(N-k,:) \\
x(k,:)
\end{bmatrix}
=
\begin{bmatrix}
\tilde{r}(N-k,:)
\end{bmatrix} \,,
\tag{4}
$$

where the matrices F and L correspond to the cross-correlation with the first and the last rows, respectively. The right-hand side $\tilde{r}(N-k)$ is obtained from the $(N-k)$-th row of the auto-correlation function $r(i,j)$ from which the contribution of already recovered rows $1, 2, \ldots, k-1$ and $M-2, M-3, \ldots, M-k+1$ has been subtracted. For a more detailed description we address the reader to the original articles. The most appealing property of this algorithm is that it requires only a small, known in advance, number of iterations until the whole signal $x(m,n)$ is recovered. The authors also provided conditions that they

---

[1]Here we use MATLAB notation, where a colon (:) denotes the entire vector of indexes of the corresponding dimension.

claimed were sufficient to guarantee uniqueness of the reconstruction. The latter, however, were proven to be incorrect (see [5]). Nonetheless, the biggest problem with this algorithm is not the non-uniqueness of reconstruction, because, as we already mentioned, the two-dimensional phase retrieval solution is usually unique (up to trivial transformations) to start with. The main difficulty that makes the algorithm impractical for all but tiny problems is its numerical instability. It can easily be shown that the error grows exponentially due to the recurrent nature of the algorithm. Even if we assume that the measurements are ideal, containing absolutely no error, each iteration of the algorithm will introduce a small error due to the finite computer precision. In the next iteration, the norm of this error will be increased by a factor proportional to the condition number of the matrix $[F \,|\, L]$. The new error will be further amplified (by the same factor) in the next iteration, and so on. This will result in extremely fast (exponential) error growth. This is demonstrated in Figure 4 where a small $(128 \times 128$ pixels) image was reconstructed by the HQ algorithm in the horizontal direction, that is, reconstructing column after column. This exponential growth is observed whenever the condition number of the matrix $[F \,|\, L]$ is greater than one. It can equal unity in some very special cases, for example, when the known boundaries contain a single delta function. This situation was considered in $[4, 5]$, and in $[2]$, although in the latter it was not used directly for the reconstruction—the authors added this condition to guarantee the uniqueness of the reconstruction. Moreover, all of them observed empirically that the reconstruction was faster when the known part contained more energy. This observation is common, even though the authors use different reconstruction methods. None of them, however, provided an explanation for this phenomenon. In the next section we present our reconstruction routine and explain why a "strong" known part leads to a fast and stable reconstruction. Additionally, we will consider the influence of noise in the measurement—another issue that has been largely overlooked in previous works despite its enormous importance.

3.   Our reconstruction method

First, we must consider the source of the known boundary. In microscopy, it is often natural to create a mask (for transparencies) or a bed (for light reflecting objects) whose boundaries are known and designed in a way that leads to easy image reconstruction. An example of such a mask is shown in Figure 1.
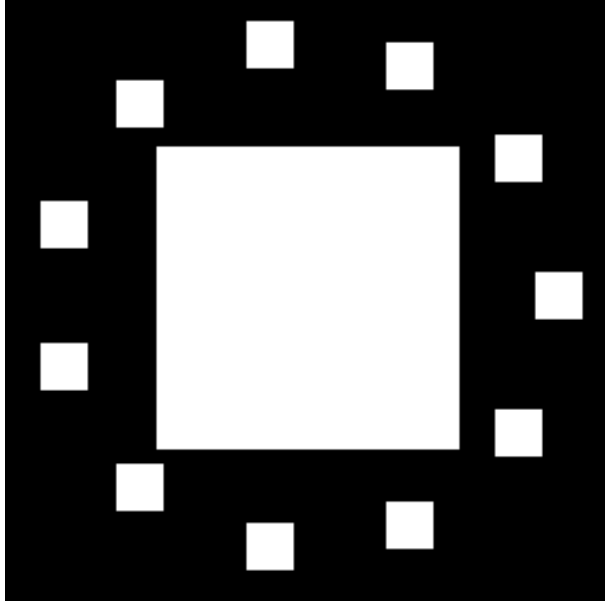
Fig. 1: Artificial mask design.

The object here is assumed to be transparent, so the white areas correspond to simple windows in an opaque material (shown in black). The big window in the center is where the object is placed. The whole construction is then illuminated by a coherent plane wave such that the small windows in the mask can be assumed to be of known intensity. There is nothing special about the plane wave here, the most important point is that part of the image is known. This setup allows us to formulate the following minimization problem to find the unknown signal $x(m, n)$:

$$\min_{x} \quad \||\mathcal{F}[x + b]| - r\|^2$$
$$\text{subject to} \quad x(m \in \mathcal{O}_m, n \in \mathcal{O}_n) = 0\,, \tag{5}$$

where r denotes the measured Fourier magnitude of the entire signal, b represents the known part (boundary) and $(\mathcal{O}_m, \mathcal{O}_n)$ designate the location of the off-support parts of x (basically, these are the locations occupied by the mask except the central window where the object is located). Note that we, again, use x to denote the reconstruction result because, in general, it may not be equal the sought signal z.

Note tat the mask can, in principle, be constructed in a way that makes the reconstruction trivial: it is sufficient to place only one infinitesimally small window (a delta function) at a sufficient distance from the object. In this case the mask would satisfy the holography conditions and the reconstruction could be as easy as applying a single Fourier transform. However, generating a delta function is not possible in practice. Practical mask design must

balance between the production costs/difficulties and how helpful it is for the reconstruction process. Hence, here we use simple square windows of relatively large size, located close to the object. Hence, a non-iterative reconstruction is not possible. However, using a quasi-Newton optimization method to solve (5) we were able to get good results, as demonstrated in Section 4.

Before we proceed to the numerical results, it is pertinent to discuss briefly the design of the mask. It is not known at the moment what is the best way to design the mask so that the reconstruction would be fast and robust. Our experience shows that the mask should have a strong "presence" in all frequencies. More precisely, we can explain why this situation leads to a fast reconstruction. Consider a single element in the Fourier space. The contribution of the known part (denoted by $\hat{b}_i$) is fully determined, that is, its magnitude and phase are known. Hence, the full Fourier domain data at this location is the sum $\hat{x}_i + \hat{b}_i$ that is located somewhere on the red circle shown in Figure 2. Note that the phase $\angle(\hat{b}_i + \hat{x}_i)$ is known up to $\pi/2$ radians when $\alpha \leq \pi/4$ (see Figure 2). That is, when

$$\arcsin\left(\frac{|\hat{x}_i|}{|\hat{b}_i|}\right) \leq \frac{\pi}{4} \iff |\hat{b}_i| \geq \sqrt{2}|\hat{x}_i|. \tag{6}$$
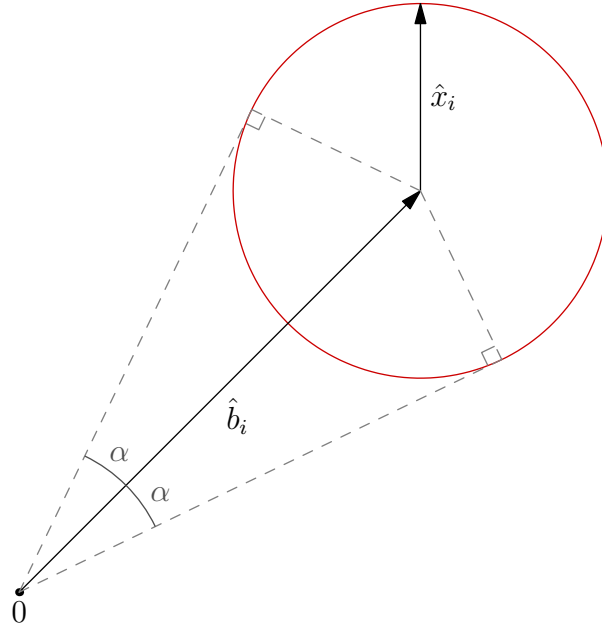

Fig. 2: "Strong" known part leads to the phase retrieval problem with approximately known Fourier phase.

The relation in Equation (6) is quite interesting: if the energy concentrated in the known

part is at least twice as large as the energy in the unknown part[2], the problem reduces to phase retrieval with Fourier phase known within the limit of $\pi/2$ radians. Hence, according to our results in [9], *any* reasonable algorithm will be able to reconstruct the unknown part. Moreover, the larger the ratio $|\hat{b}_i|/|\hat{x}_i|$, the smaller the Fourier phase uncertainty is. However, there is another aspect that is important in practice—in all physical experiments, the measurements inevitably contain some noise. In our case, where the measurements are light intensity, the noise is well approximated by the Poisson distribution. This means that stronger intensity will result in more noise (though, the signal to noise ratio (SNR) usually increases with the intensity growth). Hence, making the known part too strong compared with the sought signal will make the measurement noise more dominant with the respect to the unknown part and will eventually arrive at the level where the reconstruction is not possible. Hence, one should not just increase the intensity of the known part, as this leads to poor reconstruction quality in case of noisy measurements[3]. The most appealing approach would be to design a mask that would use the limited power in an efficient way. Without any a priori knowledge about the sought signal, it seems that the optimal way would be to create a mask whose Fourier domain power is spread evenly over all frequencies. Note the relation to the holographic reconstruction where a single small window (delta function) is used, because the Fourier domain power of a delta function is the same over all frequencies. Instead of using a delta function we can obtain a very good approximation to the uniform power spectrum if some randomness is added to the mask windows. It can be random shape of the windows or random values/phases across them. Making random shapes may be more difficult than adding a diffuser into a square window. Hence, in Section 4 we demonstrate the reconstruction with mask windows of constant intensity and random intensity. As expected, adding randomness improves the reconstruction speed and noise-immunity.

4.   Numerical results

We experimented with several images, however, the results provided below are limited to two images of size 128×128 pixels. These were chosen to represent two different classes of objects. One is a natural image "Lena" already used in our previous experiments. Another is the Shepp-Logan phantom. These images are very popular in other fields. "Lena" is a classical benchmark in the image processing community, because it has a lot of features and delicate details. The second image, "phantom", is often used as a benchmark in MRI related algorithms. However, its piece-wise constant nature can approximate well objects that are often investigated in microscopy, for example, cells. Besides the above differences these two images differ by their support. Lena's support is tight, that is, it occupies all the space and

---

[2]The energy must also be distributed "well" in the frequency domain.

[3]Here we assume that the noise grows with the signal, as indeed happens with Poissonian noise.

no shifts are possible. The phantom, on the other hand, has non-tight support. As we saw earlier, this is an important property for phase retrieval algorithms—complex valued images with non-tight support are much more difficult for the current reconstruction methods like HIO. The image intensities (squared magnitude) is shown in Figure 3.

<center>(a)   (b)</center>

Fig. 3: Test images (intensity).

The images were tested in two different scenarios: first, they were assumed to be real-valued and non-negative; second, they were made complex-valued by adding a phase distribution. Here we present the results for the case where the objects' phases were chosen to be proportional to their intensity (scaled to the interval $[0, 2\pi]$). This choice corresponds to the case where the phase changes in a relatively smooth manner. However, all our experiments indicate that the particular phase distribution has little effect on the reconstruction, except the case where the object can be assumed real non-negative. In this case the non-negativity prior can be used to speed-up the reconstruction and improve its quality in the case of noisy measurements, as is evident from Figures 6, and 7.

Let us begin with a demonstration of the HQ algorithm results. Following our discussion in Section 2 we expect the error to grow exponentially as we progress from the boundaries toward the image center. This expectation is confirmed by the results shown in Figure 4. Note that the reconstructed intensity in Figure 4b cannot convey the true error because the image storage format clips all values outside the interval $[0, 1]$. The true error is evident from Figure 4c.
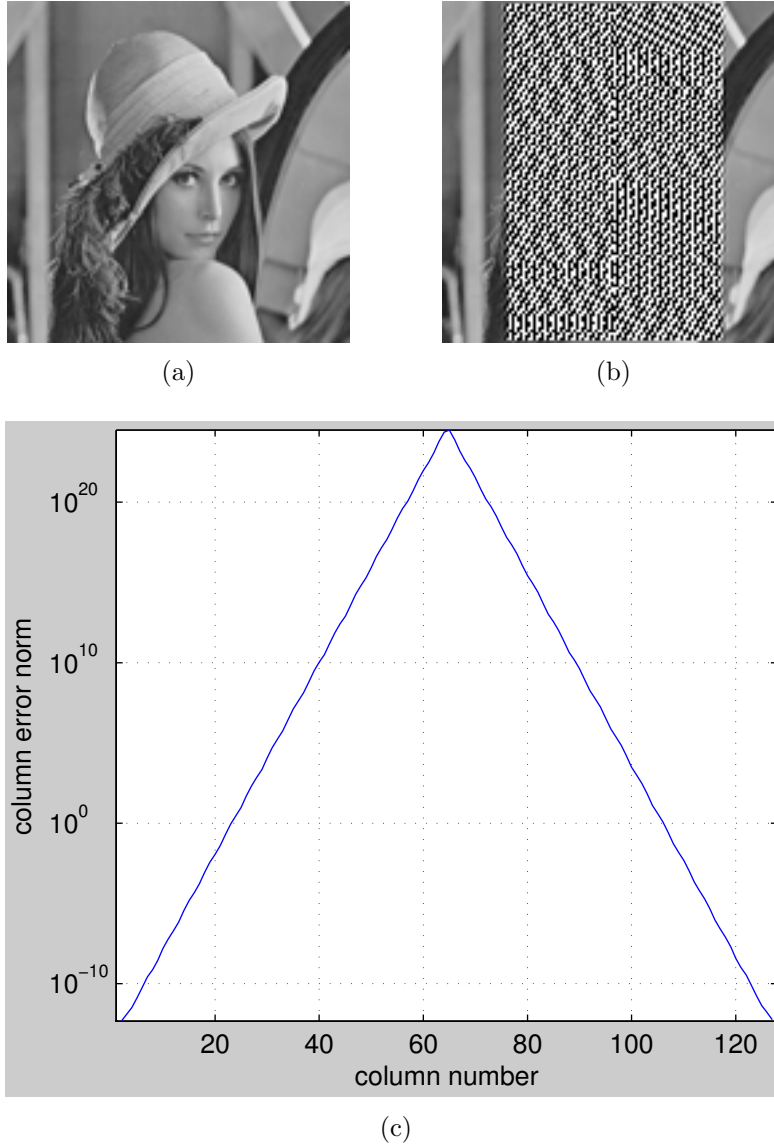
<center>8</center>

Fig. 4: Error growth in the Hayes-Quatieri recursive algorithm: (a) original image, (b) reconstructed by the HQ algorithm, and (c) actual error growth in the HQ algorithm.

In theory, we can use the same boundary conditions as the HQ algorithm, however, our experiments indicate that the optimization routine used in our method is prone to stagnation when the known boundary (or image part in general) carries little energy. This is, of course, in agreement with our discussion in the previous section: when the boundary caries little energy it does not provide enough information about the Fourier phase. This, in turn, causes line-search optimization algorithms to stagnate (see [9]). The stagnation can be alleviated by

9

designing a mask that makes the reconstruction much easier. In our approach the emphasis is also put on the simplicity of the mask design. Hence, we used the simple mask shown in Figure 1. The mask is of size 256×256 pixels with 11 square windows of size 20×20 located approximately on a circle of radius 100 pixels (see Figure 1). The objects are placed in the middle of the mask where a special window of size 128×128 pixels is provided. The area outside this special window comprises the known boundary. When placed in the mask, the test images look as shown in Figure 5.
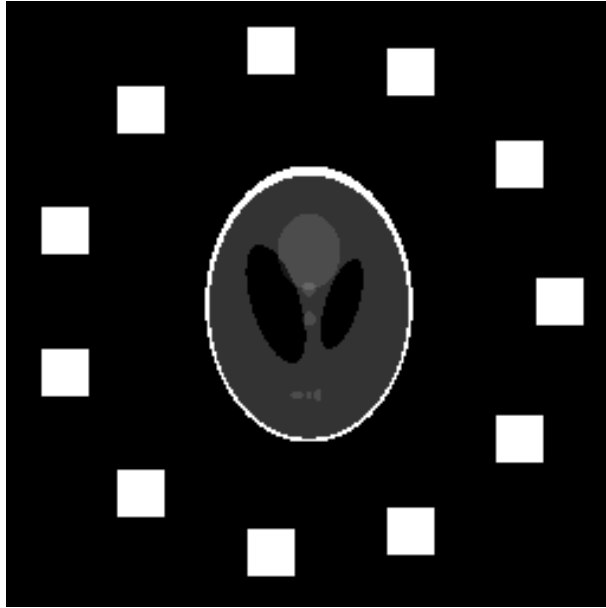
Fig. 5: Artificial mask with test images.

Note also that the algorithm based on the minimization problem defined in Equation (5) is very naive and does not try to use the "approximately known Fourier phase" in the case were the energy in the known part is sufficiently large. However, our main goal is to show the influence of the mask design. This influence is essentially independent of the algorithm

used for reconstruction.

Let us demonstrate how the energy contained in the known part affects the reconstruction. Recall that our expectation is that a mask that has a strong "presence" in all Fourier frequencies will better suit the reconstruction process in the noise-less case. The results shown in Figures 6, and 7 fully support this conjecture. These figures present the reconstruction speed (the objective function minimization rate) for three different magnitudes of the mask's windows values: 5, 10, and 100. The lowest value (5) was chosen so as to be close to the theoretical ratio between the energy in the known and the unknown parts that is required for guaranteed reconstruction, as defined in Equation (6). However, as we can see the value of 5, and even the value of 10 did not result in perfect reconstruction. This phenomenon has two reasons: first, doubling the amount of energy in the known part compared to the unknown part is not sufficient—this energy must be distributed properly in the Fourier domain; second, it may be attributed to the simplicity of the chosen reconstruction algorithm. However, the second reason is less likely in view of the following experiment. In an attempt to create a "better" distribution of the mask's energy in the Fourier domain we added some "randomness" to the mask by modulating (multiplying) the flat values across the mask's windows with random values in the interval $[0, 1]$. This step improved the speed of the reconstruction and its robustness, as is evident from Figures 8, and 9. "Random" mask are good for signals without any estimate of the Fourier magnitude, on the other hand, when the Fourier magnitude of the sought signal is known precisely, the ideal mask would have the same magnitude multiplied by two. Hence, the ideal mask would be the sought signal multiplied by two. This is, of course, is not a practical situation. In practice, however, one may know approximately the Fourier magnitude of the sought signals, in this case the best mask can be obtained by another run of a phase retrieval method in order to match the Fourier magnitude of the sought signal.
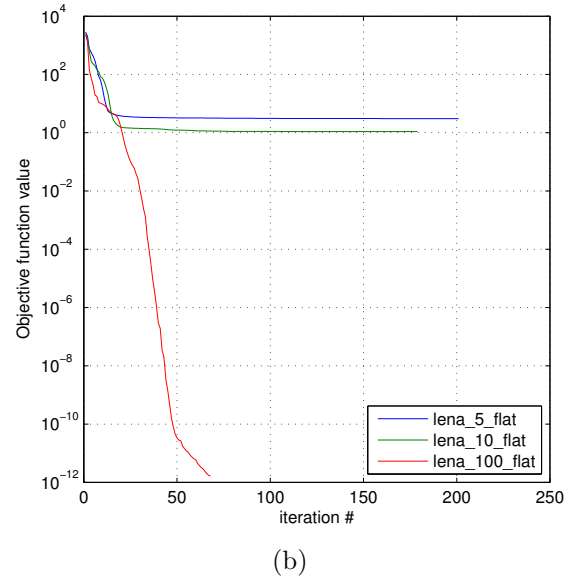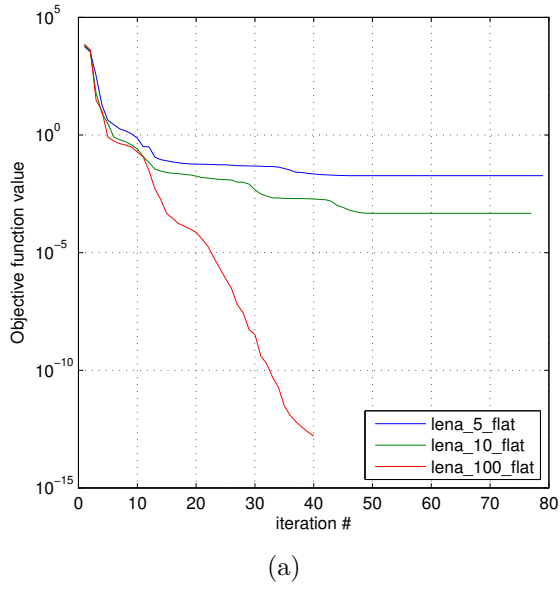
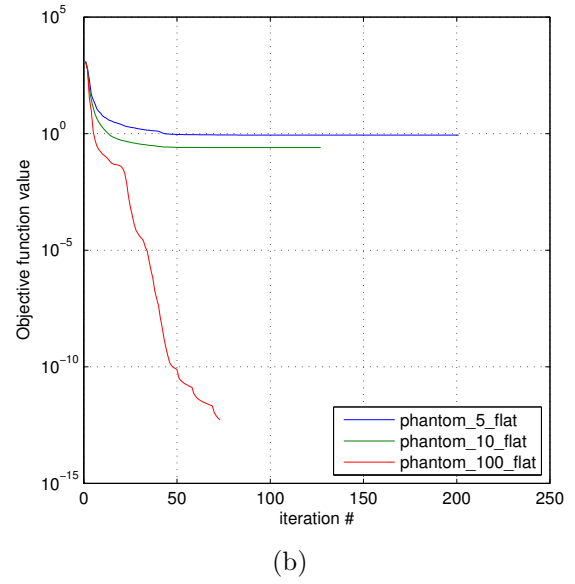Fig. 6: Lena's reconstruction speed with flat mask values: (a) real-valued, (b) complex valued.
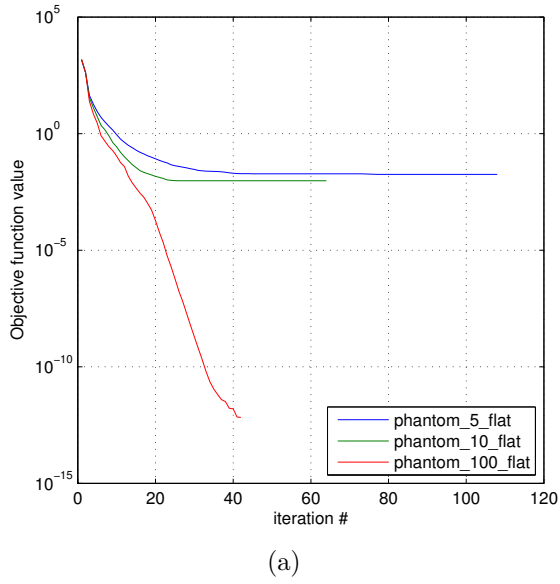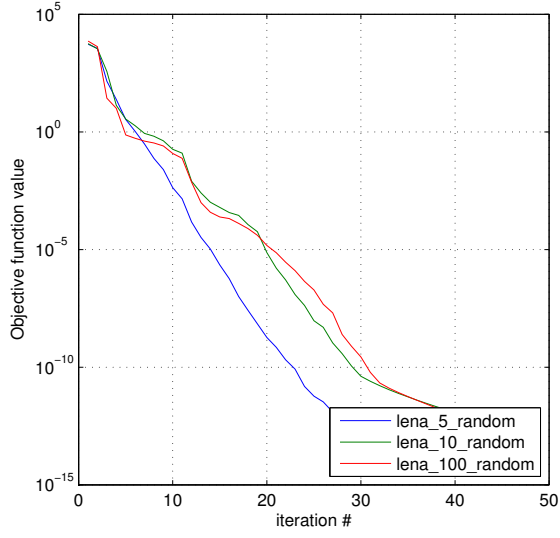

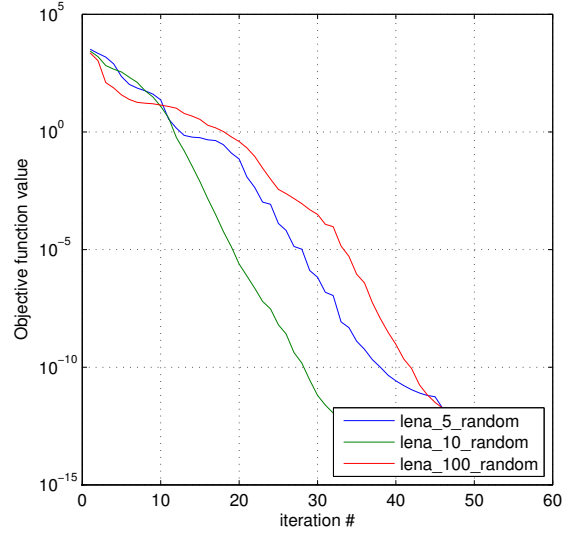Fig. 7: Phantom's reconstruction speed with flat mask values: (a) real-valued, (b) complex valued.
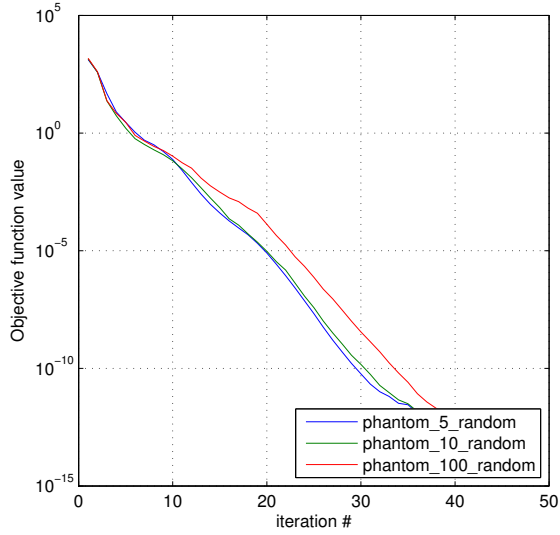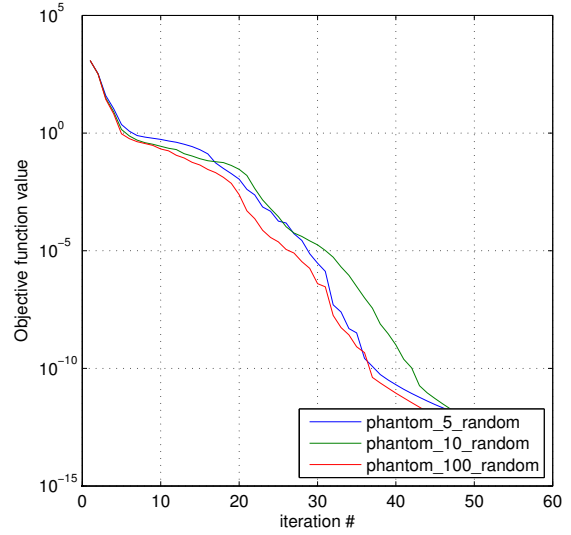
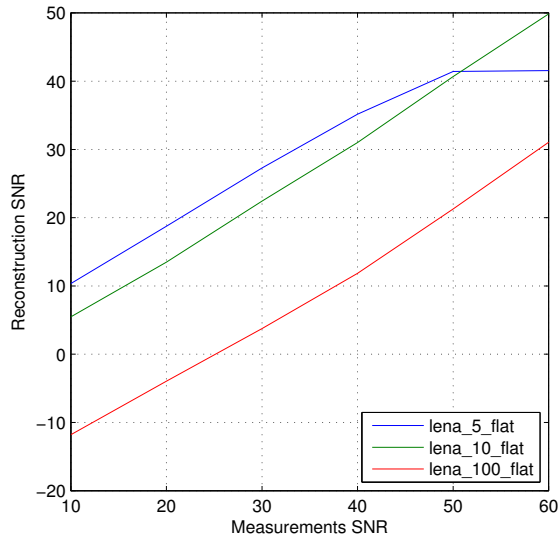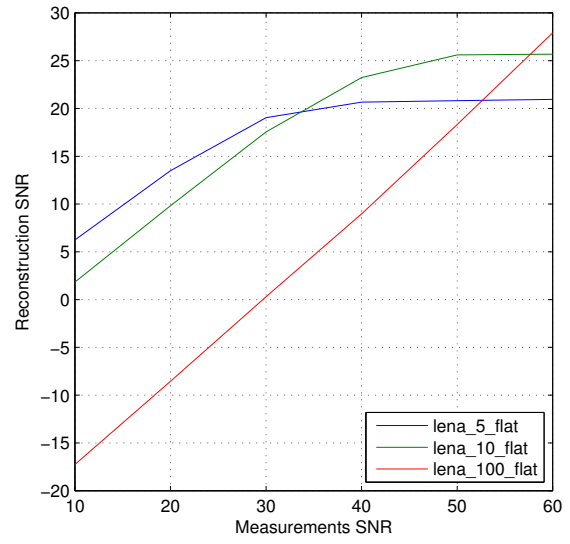Fig. 8: Lena's reconstruction speed with random mask values: (a) real-valued, (b) complex-valued.

Fig. 9: Phantom's reconstruction speed with random mask values: (a) real-valued, (b) complex-valued.

Note that now the reconstruction is successful for all mask values and is very fast. However, despite this fast convergence, one must be careful not to put too much energy into the known part. This approach may harm the reconstruction quality of the unknown part when there is noise in the measurements, as we demonstrate next. In these experiments the measurements (intensity values) were contaminated with Poisson noise with different SNR ranging from 10 to 60 decibels. As is evident from Figures 10 and 11, the more energy is concentrated in the known part the worse is the reconstruction quality of the unknown part. This phenomenon is, of course, expected. The Poisson noise is signal dependent: higher intensity results in more noise. However, the intensity (energy) of the unknown part remains constant, hence the noise becomes more and more significant compared to it. To obtain the best result one would like to design a mask whose power spectrum will correlate well with the power spectrum of the sought signal. Unfortunately, this approach cannot be implemented, because designing such a mask requires a priori knowledge of the sought signal's Fourier magnitude, which is unavailable in our case. However, it may be a good approach when the Fourier magnitude of the sought signal is known *approximately*.
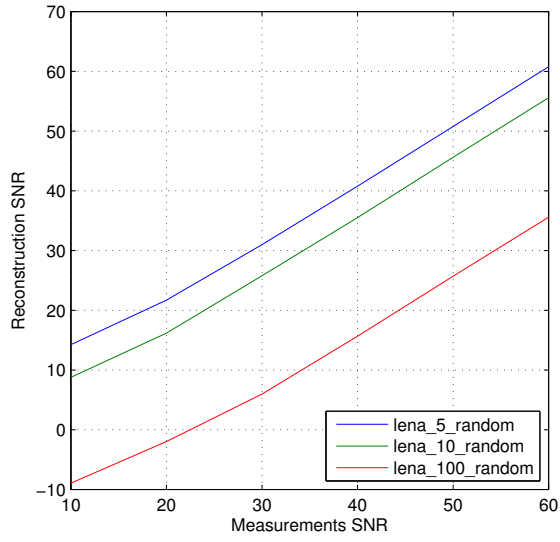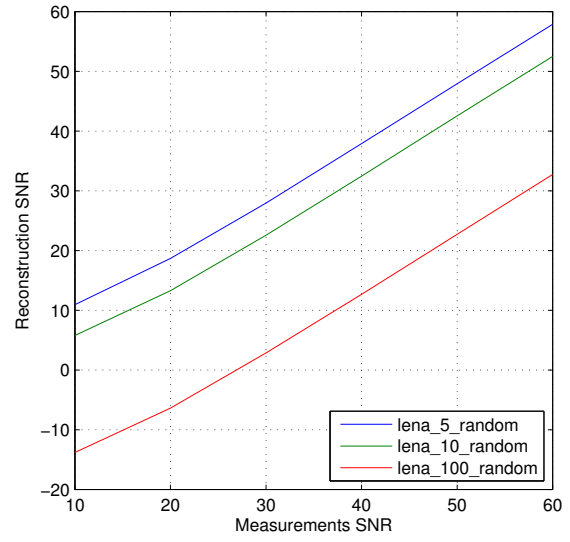
Fig. 10: Lena's reconstruction quality: (a) real-valued with flat mask value, (b) complex-valued with flat mask value, (c) real-valued with random mask values, (d) complex-valued with random mask values.
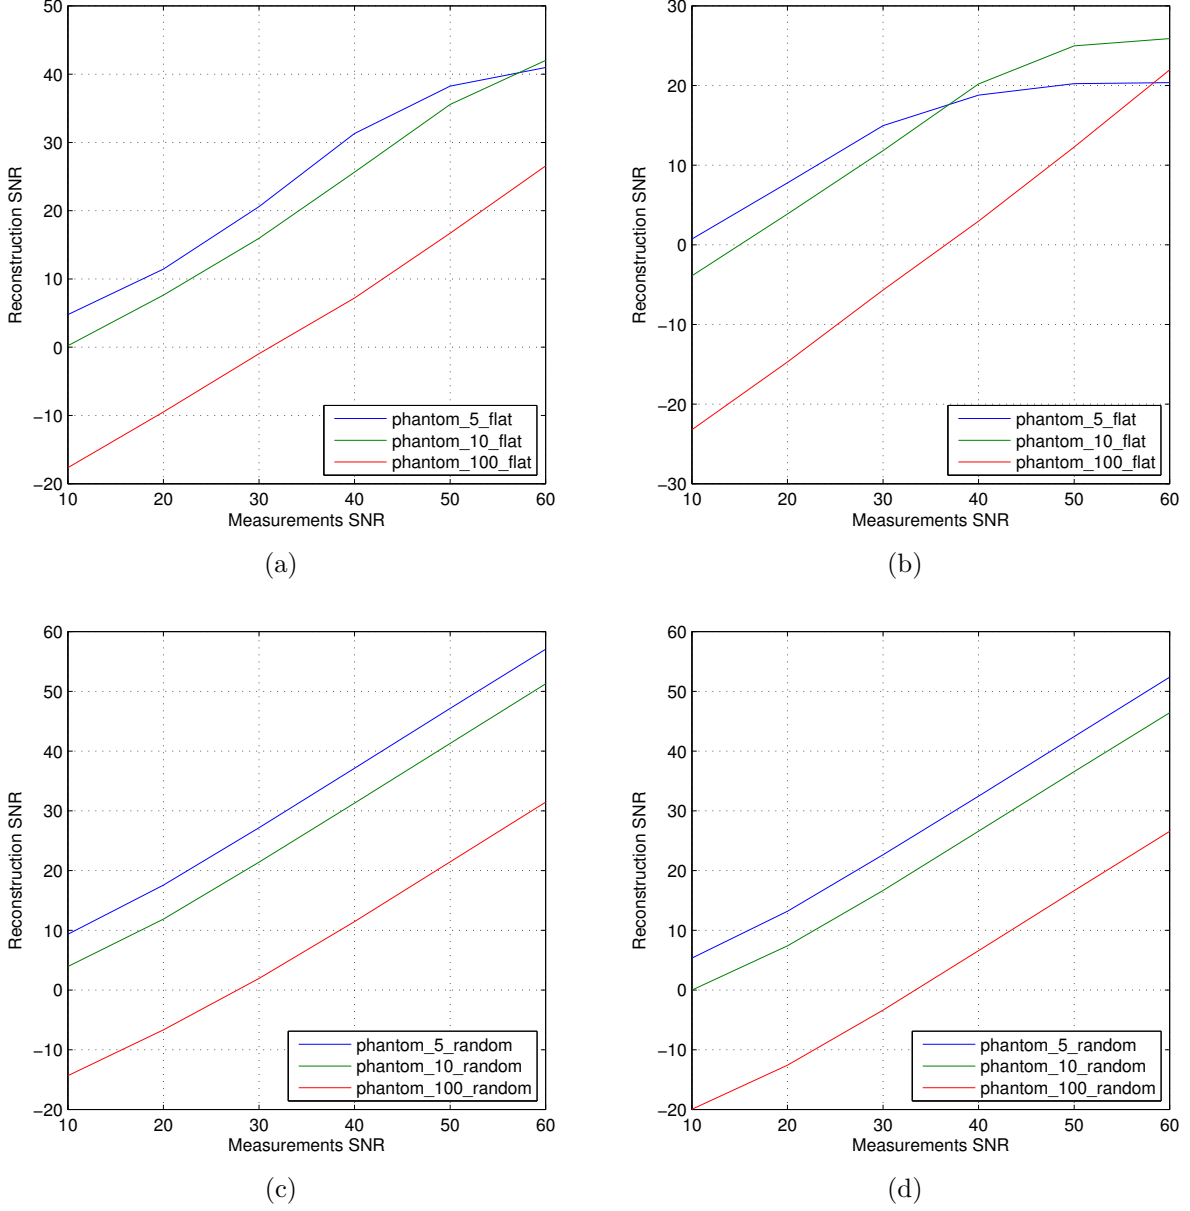
Fig. 11: Phantom's reconstruction quality: (a) real-valued with flat mask value, (b) complex-valued with flat mask value, (c) real-valued with random mask values, (d) complex-valued with random mask values.

5. Concluding remarks

In this chapter we considered the problem often met in the Fourier domain holography: signal reconstruction form the Fourier magnitude of the sum of the sought signal and a reference

beam. We provided an explanation to the fact observed in practice: why a strong reference beam leads to a faster reconstruction for a variety of reconstruction methods. Based on this explanation we suggested a "good" boundary (reference beam) design. The latter problem (reference beam design) requires more research, as the optimal reference beam must satisfy at least two requirements in the presence of signal dependent noise. For example, in the case of Poissonian noise (or any other noise model in which the noise level grows with the signal intensity) the optimal reference beam must be simultaneously "strong" (to aid the reconstruction process) and "weak" (to alleviate) the destructive influence of the noise.

In general, when the Fourier magnitude of the sought image is known approximately, the best mask should have the power spectrum that is about two times larger than that of the sought signal (in each frequency). If the power spectrum of the sought images is unknown, the mask should have a strong presence in all frequencies. In this case, it seems that the best design would be based on some "randomness" in the mask values or shapes.

References

1. Yu. M. Bruck and L. G. Sodin. On the ambiguity of the image reconstruction problem. *Optics Communications*, 30(3):304–308, September 1979.
2. M. A. Fiddy, B. J. Brames, and J. C. Dainty. Enforcing irreducibility for phase retrieval in two dimensions. *Optics Letters*, 8(2):96–98, February 1983.
3. J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982.
4. J. R. Fienup. Reconstruction of objects having latent reference points. *Journal of the Optical Society of America*, 73(11):1421–1426, November 1983.
5. J. R. Fienup. Phase retrieval using boundary conditions. *Journal of the Optical Society of America A*, 3(2):284–288, February 1986.
6. M. Hayes. The reconstruction of a multidimensional sequence from the phase or magnitude of its fourier transform. *Acoustics, Speech, and Signal Processing [see also IEEE Transactions on Signal Processing], IEEE Transactions on*, 30(2):140–154, 1982.
7. M. Hayes and T. Quatieri. The importance of boundary conditions in the phase retrieval problem. In *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '82.*, volume 7, pages 1545–1548, 1982.
8. Monson H. Hayes and Thomas F. Quatieri. Recursive phase retrieval using boundary conditions. *Journal of the Optical Society of America*, 73(11):1427–1433, November 1983.
9. Eliyahu Osherovich, Michael Zibulevsky, and Irad Yavneh. Approximate Fourier phase information in the phase retrieval problem: what it gives and how to use it. *Journal of the Optical Society of America A*, 28(10):2124–2131, October 2011.